Welfare of Sequential Allocation Mechanisms for Indivisible Goods

Haris Aziz
Data61 and UNSW,
Sydney 2033, Australia

Thomas Kalinowski
University of Newcastle
Newcastle, Australia

Toby Walsh
UNSW and Data61,
Sydney 2033, Australia

Lirong Xia
RPI
NY 12180, USA

arXiv:1511.08280v1 [cs.AI] 26 Nov 2015

Abstract

Sequential allocation is a simple and attractive mechanism for the allocation of indivisible goods. Agents take turns, according to a policy, to pick items. Sequential allocation is guaranteed to return an allocation which is efficient but may not have an optimal social welfare. We consider therefore the relation between welfare and efficiency. We study the (computational) questions of what welfare is possible or necessary depending on the choice of policy. We also consider a novel control problem in which the chair chooses a policy to improve social welfare.

Introduction

Due to economical, environmental and political concerns, we often want to do more with fewer resources and to do so more fairly. One way to achieve this is to use computing power to improve the efficiency and equitability of the allocation. One important and challenging case is the fair division of indivisible goods. This captures a wide range of problems including allocating classes to students, landing slots to airlines, players to teams, and houses to people.

A simple but popular mechanism to allocate indivisible goods is *sequential allocation* (Brams and Taylor 1996). Agents simply take turns to pick items. The sequential allocation mechanism leaves open the particular order used to take turns (the so called "policy"). Is it fairest perhaps to have a balanced alternating policy in which items are allocated in rounds, each agent picks one item in each round, but we reverse the order of the agents after each round. Indeed, there are real world settings like course allocation at the Harvard Business School where the policy is chosen at random from a space of balanced alternating policies as a means of ensuring (procedural) fairness.

The choice of policy impacts the welfare of the resulting allocation. This raises the question of what social welfare can or must be achieved. Do we *necessarily* achieve a minimum acceptable welfare whatever policy is chosen? Is is *possible* that the welfare is above a required minimum? What is the *maximum* welfare that can be achieved? What is the *minimum* welfare that will be achieved? These questions are closely related to an interesting control problem. Can a (benevolent) chair choose a policy to improve or maximize welfare? These questions are also related to the expected welfare when the policy is chosen at random. For example, the expected welfare is between the minimum welfare that is necessary and the maximum welfare that is possible (strictly so when they are different).

We are not the first to consider what allocations are possible or necessary depending on the choice of policy. Aziz, Walsh and Xia (2015) ask what item or (sub)set of items can or will be allocated depending on the choice of policy. By comparison, we consider here not the items allocated but the welfare achieved. We are also not the first to consider the policy that maximizes welfare. For example, Bouveret and Lang (2011) ask what is the "optimal" policy that maximizes the *expected* egalitarian or utilitarian social welfare. However, their model supposes that the ordinal preferences are not known by the chair and optimality is in expectation under two extreme assumptions: full independence (all rankings equiprobable) and full correlation (identical rankings). By comparison, we consider here the case where the we know the *exact* utilities and can therefore maximize the *actual* social welfare.

Sequential allocation is an ordinal mechanism (it merely requires agents to declare an ordering over items). In many of our settings, however, we suppose that we know the agents' utilities. This may be because we know the ordinal preferences but utilities can be easily computed from these (e.g. Borda or lexicographical utilities). In other cases, we might suppose that we have elicited the agents' general utilities, and we then compute a policy to maximize welfare which we announce and use to allocate items using sequential allocation. Even in this more complex setting, we retain some of the advantages of a purely ordinal mechanism. For instance, it is easy for the agents to verify that the policy is fair (e.g. it is a balanced alternating policy), and the part of the mechanism allocating items according to the declared policy has been applied correctly.

Welfare and efficiency

We first consider the precise relationship between social welfare and efficiency. We suppose that there are n agents being allocated m items. Agents have additive utilities over the items. Agents convert these into a strict ordinal ranking over items, breaking any ties in utility in some fixed way. The welfare of an agent is simply the sum of the utilities of the items allocated to that agent. The utilitarian welfare is the

sum of the welfare of the agents, whilst the egalitarian welfare is that of the worst off agent (or agents). The sequential allocation mechanism is parameterized by the policy, the order in which agents pick items. For example, with the policy 123321, agent 1 picks first, then agent 2, then agent 3 before we repeat in reverse. An allocation is an assignment of items to agents. One allocation *Pareto improves* another iff each agent has at least the same utility in the first, and there is at least one agent where the utility is greater. An allocation is *Pareto efficient* iff there is no allocation which Pareto improves it. For every Pareto efficient allocation, there exists a policy such that sincere picking with this policy generates this allocation. We can construct this policy using the greedy algorithm in the proof of Proposition 1 in (Brams and King 2005). The reverse, however, is not true. Sincere picking may not return a Pareto efficient allocation.

Remark 1. *Sincere picking can generate allocations that are not Pareto efficient.*

Proof: Consider the policy 1221. Suppose the agents' utilitiea are as follows

	a	b	c	d
1	5	4	2	0
2	8	2	1	0

Both agents have the same total utility over the items. Sincere picking gives items a and d to agent 1 and items b and c to agent 2. This gives an utility of 5 to agent 1 and of 3 to agent 2. If they swap allocations, then the utility of agent 1 increases to 6, and of agent 2 to 8. Hence, sincere picking leads to an allocation that is not Pareto efficient, and does not have the optimal egalitarian or utilitarian social welfare. ◇

We contrast this observation with Proposition 1 in (Brams and King 2005). This looks just at the rank of items in an agent's preference ordering, ignoring their precise utilities. Given two sets of items S and S' with $|S| = |S'|$, an allocation of items S to an agent *dominates* the allocation of items S' iff for every item in $S - S'$ there is a different item in $S' - S$ that is strictly less preferred. They then define an ordering, *ordinal efficiency* in terms of such domination. This is a strictly weaker ordering than Pareto efficiency which is defined in terms of utilities rather than ordinal rankings.

Proposition 1 in (Brams and King 2005) demonstrates that ordinal efficiency corresponds *exactly* to allocations generated by sequential allocation supposing sincere picking. On the other hand, only a subset of the allocations returned by sequential allocation are Pareto efficient. And only a subset of these maximize the egalitarian social welfare. We highlight the fact that there exists an allocation with the maximum possible egalitarian welfare.

Remark 2. *There exists an allocation with the maximum possible egalitarian social welfare that is also Pareto efficient.*

The argument is as follows. Among all allocations with maximum egalitarian welfare choose one with the largest utilitarian welfare. This allocation is clearly Pareto efficient.

It follows quickly that there always exists a policy for sequential allocation that gives an allocation with the maximum possible egalitarian social welfare supposing sincere picking. Note that the proof does not rule out other allocations which maximize egalitarian social welfare which are not ordinal efficient, and which cannot be generated by sequential allocation with sincere picking.

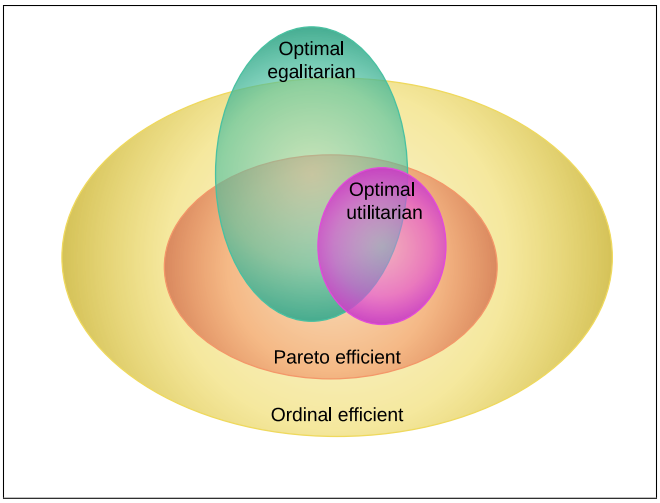

Figure 1: Relationship between different allocations. Sequential allocation with sincere picking can generate any ordinal efficient allocation.

Example 1. *Suppose we have three agents (1 to 3), three items (a to c), and Borda utilities. Let agent 1 have a preference order bac, agent 2 have abc, and agent 3 have acb. Then the allocation which gives a to agent 1, b to agent 2 and c to agent 3 maximizes the egalitarian social welfare. However, there is no policy for sequential allocation that will return such an allocation supposing agents pick sincerely as no agent gets a first choice item.*

Maximizing the utilitarian social welfare also does not conflict with Pareto efficiency. In this case, we point out the well-known fact that *any* allocation that maximizes utilitarian social welfare is Pareto efficient.

Remark 3. *Any allocation with the maximum possible utilitarian social welfare is also Pareto efficient.*

The argument is as follows. Consider any allocation that has the maximum possible utilitarian social welfare. Suppose there exists another allocation which Pareto improves it. Then the utility of every agent does not decrease. This means that the sum of their utilities must increase. This contradicts the assumption that we have the maximum possible utilitarian social welfare.

Again it follows quickly that there exists a policy that gives an allocation with the maximum possible utilitarian social welfare supposing sincere picking.

Possible and necessary welfare

Since sequential allocation may not return allocations that are optimal from either an egalitarian or utilitarian perspective, we turn to the (computational) questions of what social welfare is possible or necessary. Note that throughout this paper, we suppose agents pick sincerely. Whilst strategic behaviour may be beneficial, risk averse agents will tend

to pick sincerely, especially when the policy and/or utilities are private information. Nevertheless, it would be interesting future work to consider agents acting strategically (Kalinowski et al. 2013).

POSSIBLE/NECESSARY UTILITARIAN/ EGALITARIAN WELFARE
Input: a set of n items, m agents each with utilities over the items, a class of policies, and an integer t.
Question: Is there a policy/Does every policy result in an allocation with an utilitarian/egalitarian social welfare of t or greater supposing agents pick items sincerely?

The possible welfare questions answer a *policy control* problem: can the chair choose a policy to achieve a given social welfare? Similar control problems have been considered previously (Aziz, Walsh, and Xia 2015) but with the goal of allocating particular items to agents, rather than, as here, of achieving a particular welfare. Note that we suppose we know the (private) utilities of the agents. Our complexity results can be seen as lower bounding the computational complexity when we only have partial information about the actual utilities. Alternatively, we may relax the assumption that we know the actual utilities. For example, as in (Brams, Edelman, and Fishburn 2003; Bouveret and Lang 2011; Kalinowski, Narodytska, and Walsh 2013; Baumeister et al. 2014), we might suppose that the utilities are simple functions of the ordinal rank (e.g. Borda, lexicographical or quasi-indifferent scores). As this is a special case of general utilities, any result that control takes polynomial time in the general case will map onto a polynomial time result in this more restricted setting.

When we prove that a particular possible or necessary welfare problem takes polynomial time to solve, we will typically do so by answering a closely related maximization or minimization problem. Such problems are interesting in their own right.

MAXIMUM/MINIMUM UTILITARIAN/ EGALITARIAN WELFARE
Input: a set of n items, m agents each with utilities over the items, and a class of policies.
Output: The maximum/minimum utilitarian/egalitarian social welfare possible over all policies supposing agents pick items sincerely.

All possible policies

If any policy is possible, it is easy to maximize the utilitarian social welfare. The chair just need to choose a policy that gives items to the agents which value them most.

Theorem 1. *The* MAXIMUM *and* POSSIBLE UTILITARIAN WELFARE *problems are polynomial time solvable.*

Proof: We order the items by the maximum utility assigned by any agent. Ties can be broken in any way. We then construct the policy that allocates items in this order choosing the agent who gives an item the greater utility. No allocation can do better than this. $\diamond$

The MINIMUM EGALITARIAN WELFARE problem also takes polynomial time to solve. It is always zero. On the other hand, the POSSIBLE EGALITARIAN WELFARE problem is intractable in general, even in the special case that all the agents have identical utilities for the items.

Theorem 2. *The* POSSIBLE EGALITARIAN WELFARE *problem for m items and n agents is strongly NP-complete when $m \geqslant 2n$.*

Proof: Membership in NP is shown by giving the policy. The proof uses a reduction from numerical 3-dimensional matching. Given an integer t and 3 multisets $X = \{x_1, \ldots, x_n\}$, $Y = \{y_1, \ldots, y_n\}$ and $Z = \{z_1, \ldots, z_n\}$ of integers with $\sum_{i=1}^{n}(x_i + y_i + z_i) = nt$, this problem asks if there are permutations σ and π such that $x_i + y_{\sigma(i)} + z_{\pi(i)} = t$ for all $i \in [n]$. We construct an allocation problem over n agents and $m \geqslant 2n$ items as follows. Let $u = 1 + \sum_{i=1}^{n} z_i$. For every $j \in [n]$, there is a "big" item with utility $u + x_i + y_j$ for agent i ($i = 1, \ldots, n$) and a "small" item which all agents give utility z_j. Finally, there are $m - 2n$ items with zero utility for all agents. We ask if we can achieve an egalitarian welfare of $u + t$. To achieve this, each agent must get precisely a utility of $u + t$. This is only possible if each agent gets one big item and one small item, and $x_i + y_{\sigma(i)} + z_{\pi(i)} = t$ where $\sigma(i)$ and $\pi(i)$ denote are the indices of the big and the small item obtained by agent i. Therefore, we can achieve the egalitarian welfare of $u + t$ iff there is a solution of the original numerical 3-dimensional matching problem. $\diamond$

The same reduction proves that the MAXIMUM EGALITARIAN WELFARE problem is NP-hard to compute. In the more restricted setting that utilities are Borda scores but agents have different ordinal preferences, the POSSIBLE and MAXIMUM EGALITARIAN WELFARE problems remain NP-hard. This follows from the reduction used to prove Theorem 3 in (Baumeister et al. 2013) (due to an anonymous reviewer of the paper). It proves that deciding if there is an allocation with an egalitarian social welfare greater than or equal to some constant t is NP-complete even when utilities are Borda scores. It is easy to show that there is a policy that finds the precise allocation constructed in this reduction.

Balanced policies

It might be considered unfair to use any policy, even one in which one agent gets many more items than another. Whilst looking for allocations that maximize fairness and efficiency, Brams and King (2005) observe that *"the symbolic value of giving players equal numbers of items, such as landing slots at an airport, may be important"*. We therefore consider the restricted class of balanced policies. In a *balanced* policy, each agent gets the same number of items. For simplicity, we suppose the number of items is an integer multiple of the number of agents and add dummy items of no utility otherwise. Limiting sequential allocation to balanced policies impacts the social welfare that can be obtained.

To maximize utilitarian welfare, we cannot simply give items to the agents that value them most. This may violate balance. Despite this restriction, we can still find the policy that maximizes the utilitarian welfare in polynomial time.

Theorem 3. *The* MAXIMUM *and* POSSIBLE UTILITARIAN WELFARE *problems for balanced policies take polynomial*

time to solve.

Proof: We suppose that there are kn items to divide between the n agents. We set up a min cost max flow problem. We connect the source node to nodes representing the agents, each with a capacity of k and no cost. We connect the nodes representing agents to nodes representing the items. Each edge has a capacity of 1, and a cost equal to minus the utility that the agent assigns to the item. Finally we connect the nodes representing the items to the target node, each with an edge of capacity 1 and zero cost. We find a Pareto efficient allocation from any such flow using the top trading cycle algorithm (Shapley and Scarf 1974). A policy can be constructed that achieves this Pareto efficient allocation by again exploiting Proposition 1 in (Brams and King 2005). $\diamond$

By comparison, the NECESSARY UTILITARIAN WELFARE problem is intractable for balanced policies.

Theorem 4. *The* NECESSARY UTILITARIAN WELFARE *problem for balanced policies is coNP-complete.*

Proof: We reduce from the NECESSARY ITEM problem for balanced policies which is coNP-complete even when limited to an agent's most preferred item (Aziz, Walsh, and Xia 2015). Let one agent have utility of 1 for their most preferred item, and zero utility for all others. By comparison, let the other agents all have utility 1 for every item. Then the NECESSARY ITEM problem is equivalent to asking if an utilitarian welfare of m or more is necessary. $\diamond$

It follows that the MINIMUM UTILITARIAN WELFARE problem for balanced policies is NP-hard to compute. Restricting to balanced policies also does not change the NP-hardness of the MAXIMUM and POSSIBLE EGALITARIAN WELFARE problems. This follows almost immediately from the reduction used in the proof of Theorem 2. Note that this reduction uses policies in which some agents get less than 3 items (which are not balanced). However, such unbalanced policies can be trivially ignored as they result in poor egalitarian social welfare. Note also that when a numerical 3-dimensional matching exists, the corresponding successful policy constructed in the reduction is balanced. When utilities are specified in binary, an easy reduction from the EQUI-PARTITION problem demonstrates that the POSSIBLE EGALITARIAN WELFARE problem restricted to balanced policies is NP-complete even with just two agents who have identical utilities. Finally, the NECESSARY EGALITARIAN WELFARE problem is intractable for balanced policies.

Theorem 5. *The* NECESSARY EGALITARIAN WELFARE *problem for balanced policies is coNP-complete.*

Proof: The same reduction as in the proof of Theorem 4. $\diamond$

Recursively balanced policies

Balanced policies might still be considered unfair. For example, a policy like 11112222 favours the first agent even though it is balanced, and is guaranteed to return a Pareto efficient allocation. We therefore consider an even more restrictive class: recursively balanced policies. In such a policy, items are allocated in rounds, and each agent appears once in each round. For simplicity, we again suppose that the number of items is an integer multiple of the number of agents and add dummy items of no utility otherwise. When the number of items equals the number of agents, all balanced policies are recursively balanced. For this reason, we focus on problems where the number of items exceeds the number of agents. Recursively balanced policies include the balanced alternating policy (12211221...), as well as the Thue-Morse sequence (122121121221...). With two agents, recursively balanced policies are concatenations of 12 and 21. Other simple properties of recursively balanced policies follow immediately from their definition. For example, no agent has more than two successive picks in a recursively balanced policy. Limiting sequential allocation to recursively balanced policies may further impact the social welfare that can be obtained.

There are several situations where focusing on recursively balanced policies does not hurt welfare. For example, with Borda utilities, the expected utilitarian social welfare for two agents is not impacted by limiting allocation to recursively balanced policies. The simple alternating policy which is recursively balanced is optimal in expectation (Kalinowski, Narodytska, and Walsh 2013). Similarly for Borda utilities and small n, the expected egalitarian social welfare for two agents is not impacted. We have computed the policies that maximize expected egalitarian social welfare for up to 12 items and for each n, at least one optimal policy is recursively balanced.

In general, restricting to recursively balanced policies results in it being intractable to decide if a given egalitarian or utilitarian welfare can or must be achieved.

Theorem 6. *The* POSSIBLE EGALITARIAN *and* POSSIBLE UTILITARIAN WELFARE *problems for recursively balanced policies are NP-complete, whilst the* NECESSARY EGALITARIAN *and* NECESSARY UTILITARIAN WELFARE *are coNP-complete.*

Proof: We reduce from the corresponding problem of deciding whether the top k most preferred items of an agent are possible or necessary (Aziz, Walsh, and Xia 2015). The TOP-k POSSIBLE SET problem for recursively balanced policies is NP-complete for $k \geq 3$. We reduce this to the POSSIBLE EGALITARIAN WELFARE problem as follows. Let one agent have utility of k^2 for their ith most preferred items ($i \leq k$) and zero utility for all others. By comparison, let the other agents all have utility k^3 or greater for any item. Then the TOP-k POSSIBLE SET problem is equivalent to asking if an egalitarian welfare of k^3 or more is possible. We also reduce the TOP-k POSSIBLE SET problem to the POSSIBLE UTILITARIAN WELFARE problem as follows. Let one agent have utility of mk^2 for their ith most preferred items ($i \leq k$) and zero utility for all others. By comparison, let all the other agents have utility of k or less for any item. Then the TOP-k POSSIBLE SET problem is equivalent to asking if an utilitarian welfare of mk^3 or more is possible.

The TOP-k NECESSARY SET problem is coNP-complete for recursively balanced policies. We reduce this to the NECESSARY EGALITARIAN WELFARE problem as follows. Let one agent have total utility of k^2 for their k most preferred items and zero utility for all others. By comparison, let the other agents all have utility k^3 or greater for any item. Then

the TOP-k NECESSARY SET problem is equivalent to asking if an egalitarian welfare of k^2 is necessary. We also reduce the TOP-k NECESSARY SET problem to the NECESSARY UTILITARIAN WELFARE problem as follows. Let one agent have utility of mk^2 for their ith most preferred items ($i \leq k$) and zero utility for all others. By comparison, let all the other agents have utility of k or less for any item. Then the TOP-k NECESSARY SET problem is equivalent to asking if an utilitarian welfare of mk^3 or more is necessary. $\diamond$

Even when agents have identical utilities, these problems can remain intractable.

Theorem 7. *When allocating $2n$ items between two agents, the* POSSIBLE EGALITARIAN WELFARE *problem for recursively balanced policies is NP-complete even when agents have identical utilities given in binary.*

Proof: Membership in NP is clear. For the hardness we use reduction from PARTITION: for positive integers $a_1, \ldots, a_n$ with $a_1 + \cdots + a_n = 2B$, the problem is to decide if there is a nonempty set $I \subseteq [n]$ with $\sum_{i \in I} a_i = B$. We reduce this to the POSSIBLE EGALITARIAN WELFARE problem for two agents and $2n$ items with utilities $c_1 = 2B$, $c_{2n} = 0$, and

$$c_{2k} = c_{2k+1} = c_{2k-1} - a_k \qquad \text{for } k = 1, 2, \ldots, n-1.$$

Let $C = \sum_{i=1}^{2n} c_i$ be the sum of the utilities. Note that an egalitarian welfare of $C/2$ is equivalent to both agents achieving the same utility $u_1 = u_2$. In round k, the items with utilities c_{2k-1} and c_{2k} are allocated. From $c_{2k-1} - c_{2k} = a_k$ it follows that the difference $u_1 - u_2$ between the agents' utilities increases by a_k if agent 1 starts and decreases if agent 2 starts. Let $I \subseteq [n]$ be the set of rounds in which agent 1 starts. An egalitarian social welfare of $C/2$ is achieved if and only if

$$0 = u_1 - u_2 = \sum_{k \in I} a_k - \sum_{k \in [n] \setminus I} a_k,$$

i.e., if and only if there is a perfect partition. $\diamond$

Balanced alternating policies

The final and most restricted class of policies we consider is that of balanced alternating. This is the subclass of recursively balanced policies in which each round is the reverse of the previous. When allocating students to courses at the Harvard Business School, such a policy is chosen uniformly at random from the space of all possible balanced alternating policies. This gives a form of procedural fairness.

Theorem 8. *The* POSSIBLE EGALITARIAN *and* POSSIBLE UTILITARIAN WELFARE *problems for balanced alternating policies are NP-complete, whilst the* NECESSARY EGALITARIAN *and* NECESSARY UTILITARIAN WELFARE *are coNP-complete.*

Proof: By reduction as in the proof of Theorem 6 from the corresponding TOP-k POSSIBLE or NECESSARY SET problem restricted to balanced alternating policies. The TOP-k POSSIBLE problem for balanced alternating policies is NP-complete for $k \geq 2$, whilst the TOP-k NECESSARY SET problem is coNP-complete (Aziz, Walsh, and Xia 2015). $\diamond$

It follows that it is NP-hard to compute the probability that the Harvard Business School course allocation mechanism returns an allocation with egalitarian or utilitarian welfare greater than or equal to some given value, t.

Two agents

We now consider some special cases which are more tractable. With two agents, we can find a balanced policy that maximizes the egalitarian or utilitarian welfare in polynomial time.

Theorem 9. *The* MAXIMUM EGALITARIAN *and* MAXIMUM UTILITARIAN WELFARE *problems with balanced policies can be be solved in $O(k^2n^3)$ and $O(kn^2)$ time respectively when allocating $2n$ items between two agents with utilities (that may be different) taken from $[0, k]$.*

Proof: We put the items into some (arbitrary) order and consider how each item is allocated in turn. We construct a $2n$ step dynamic program in which the ith step corresponds to the decision of where to allocate the ith item in this order. The states of this dynamic program are triples containing the number of items allocated to the first agent, the sum of the utilities of the items so far allocated to the first agent, and the sum of the utilities of the items so far allocated to the second agent. We can compute the number of items allocated to the second agent from this. As both sums are bounded in size by $2kn$, this dynamic program has $O(k^2n^3)$ states. For the maximum utilitarian welfare, the states of the dynamic program can be simpler and just need to be pairs containing the number of items allocated to the first agent, and the sum of the utilities of the items so far allocated to both agents. $\diamond$

This result generalizes to a bounded number of agents. On the other hand, when utilities are specified in binary, an easy reduction from the PARTITION problem demonstrates that the POSSIBLE EGALITARIAN WELFARE problem is NP-complete even when the two agents have identical utilities. This is almost identical to Proposition 2 in (Bouveret and Lang 2011) which shows that deciding if there is a policy that ensures a given expected egalitarian welfare is NP-complete when the utilities of the two agents are identical.

With recursively balanced policies, we consider the case where agents have the same ordinal ranking over items.

Theorem 10. *The* MAXIMUM *and* POSSIBLE EGALITARIAN WELFARE *problems for recursively balanced policies can be solved in $O(k^2n^2)$, whilst the* MAXIMUM *and* POSSIBLE UTILITARIAN WELFARE *problems can be solved in just $O(kn)$ time when allocating $2n$ items between two agents when agents have the same ordering over items but possibly different utilities, and utilities are drawn from $[0, k]$.*

Proof: We construct a n step dynamic program in which each step corresponds to one round of allocating one item to each of the agents. The states of this dynamic program are pairs containing the sums of the utilities of items so far allocated to the two agents. As both sums are bounded by kn, this dynamic program has $O(k^2n^2)$ states. To compute the optimal utilitarian social welfare, we can use a simpler dynamic program where the states are just the sum of the utilities allocated to the two agents. $\diamond$

	all policies	balanced	recursively balanced	balanced alternating
POSSIBLE EGALITARIAN WELFARE	NPC	NPC	NPC	NPC
POSSIBLE UTILITARIAN WELFARE	P	P	NPC	NPC
NECESSARY EGALITARIAN WELFARE	P	coNPC	coNPC	coNPC
NECESSARY UTILITARIAN WELFARE	?	coNPC	coNPC	coNPC

Table 1: Summary of results: NPC=NP-complete, P=polynomial.

This result again generalizes to a bounded number of agents easily.

House allocation

Another more tractable case is house allocation, when we have only as many items as agents. In this case, we can solve the MAXIMUM and POSSIBLE EGALITARIAN WELFARE problems over all possible policies in polynomial time. We construct a graph between agents and items with edges for all items that have a utility greater than or equal to the desired egalitarian social welfare. The POSSIBLE EGALITARIAN WELFARE problem is solvable if we can find a perfect matching in this graph. To construct a satisfying policy, we find a Pareto efficient allocation from this matching using the top trading cycle algorithm (Shapley and Scarf 1974). A policy can be constructed that achieves this Pareto efficient allocation using Proposition 1 in (Brams and King 2005). This tractability of this case suggests an interesting open problem. We have proved that POSSIBLE EGALITARIAN WELFARE problem is NP-complete for $m = 2n$ but takes polynomial time for $m = n$. This leaves open the complexity in between.

Other related work

As mentioned earlier, Bouveret and Lang (2011) consider the case in which the utilities of items are simply functions of the ordinal rankings. They prove that any recursively balanced policy tends to an allocation giving the optimal expected egalitarian or utilitarian social welfare as the number of items grows, supposing sincere picking, utilities that are Borda scores and all ordinal rankings being equiprobable. In addition, they compute the optimal policies for maximizing the expected egalitarian or utilitarian social welfare under the same assumptions for up to 12 items. The optimal policies for two agents and an even number of items are recursively balanced. Kalinowski, Narodytska and Walsh (2013) prove that the alternating policy maximizes the expected utilitarian social welfare under these same assumptions. We again note that such results are about maximizing the *expected* welfare supposing limited knowledge about the utilities, whilst the results here about maximizing the *exact* welfare supposing the chair knows the *actual* utilities.

There has been some study of strategic behaviour of agents (as opposed to the chair) in the sequential allocation mechanism. It can, for example, be viewed as a repeated game. When all agents have complete information, we can compute the subgame perfect Nash equilibrium. This is unique and takes polynomial time to compute for two agents (Kohler and Chandrasekaran 1971; Kalinowski et al. 2013), but for an arbitrary number of agents, there can be an exponential number of equilibria and computing even one is PSPACE-hard (Kalinowski et al. 2013). More recently, Bouveret and Lang (2014) consider how an agent or coalition of agents can strategically mis-report their preferences in a sequential allocation mechanism supposing the other agents act sincerely. They show that the loss of social welfare caused by such manipulation is not great. For example, with Borda scoring, two agents, and the alternating policy, there was at most a 33% loss in the utilitarian welfare.

More recently, a family of rules for dividing indivisible goods among agents has been proposed that take as input the agents' ordinal rankings over the items, a scoring vector, and a social welfare aggregation function (Baumeister et al. 2013; Baumeister et al. 2014). They return the allocation that maximizes the social welfare according to this scoring rule and aggregation function. Whilst such rules have a number of desirable properties like monotonicity, they have a number of less desirable properties including a high computational complexity to compute the actual allocation (unless we have a bounded number of agents in which case we can typically use dynamic programming). This contrasts with sequential allocation where computing the allocation take just linear time. Baumeister *et al.* (2014) also compute the multiplicative/additive "price of elicitation-freeness", the worst-case ratio/difference in social welfare between such allocations and the allocation returned by sequential allocation. Whilst their results are limited to simple alternating policies, the prices are typically not great. For example, the optimal utilitarian social welfare with Borda scores is at most twice that returned by sequential allocation using simple alternation.

Conclusions

We have considered the implications on social welfare of choosing different policies when using a sequential mechanism to allocate indivisible goods. In particular, we consider the (computational) questions of what welfare is possible or necessary. The former is related to the control problem in which a (benevolent) chair chooses a policy for the sequential allocation mechanism to improve the social welfare. These questions are also related to the expected welfare when we choose a policy uniformly at random. Our results are summarized in Table 1. There are many interesting open questions. For example, how difficult is it to find a recursively balanced policy that returns a Pareto efficient allocation supposing agents pick sincerely?

References

[2015] Aziz, H.; Walsh, T.; and Xia, L. 2015. Possible and necessary allocations via sequential mechanisms. In *Proceedings of the 29th AAAI Conference on AI*. Association for Advancement of Artificial Intelligence.

[2013] Baumeister, D.; Bouveret, S.; Lang, J.; Nguyen, T.; Rothe, J.; and Saffidine, A. 2013. Positional scoring rules for the allocation of indivisible goods. In *11th European Workshop on Multi-Agent Systems (EUMAS)*, 1–14.

[2014] Baumeister, D.; Bouveret, S.; Lang, J.; Nguyen, N.; Nguyen, T.; Rothe, J.; and Saffidine, A. 2014. Axiomatic and computational aspects of scoring allocation rules for indivisible goods. In *5th International Workshop on Computational Social Choice (COMSOC)*, 1–22.

[2011] Bouveret, S., and Lang, J. 2011. A general elicitation-free protocol for allocating indivisible goods. In Walsh, T., ed., *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011)*, 73–78. IJCAI/AAAI.

[2014] Bouveret, S., and Lang, J. 2014. Manipulating picking sequences. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI'14)*. IOS Press.

[2005] Brams, S., and King, D. 2005. Efficient fair division: Help the worst off or avoid envy? *Rationality and Society* 17(4):387–421.

[1996] Brams, S., and Taylor, A. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge: Cambridge University Press.

[2003] Brams, S.; Edelman, P.; and Fishburn, P. 2003. Fair division of indivisible items. *Theory and Decision* 55(2):147–180.

[2013] Kalinowski, T.; Narodytska, N.; Walsh, T.; and Xia, L. 2013. Strategic behavior when allocating indivisible goods sequentially. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI 2013)*. AAAI Press.

[2013] Kalinowski, T.; Narodytska, N.; and Walsh, T. 2013. A social welfare optimal sequential allocation procedure. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI-2013)*. International Joint Conference on Artificial Intelligence.

[1971] Kohler, D., and Chandrasekaran, R. 1971. A Class of Sequential Games. *Operations Research* 19:270–277.

[1974] Shapley, L., and Scarf, H. 1974. On cores and indivisibility. *Journal of Mathematical Economics* 1(1):23 – 37.